

Identifying Signatures of Selection in Genetic Time Series

Alison Feder^{1,*}, Sergey Kryazhimskiy^{2,*}, Joshua B. Plotkin¹

¹Department of Biology, University of Pennsylvania, Philadelphia, PA

²Department of Organismic and Evolutionary Biology and FAS Center for Systems Biology, Harvard University, Cambridge, MA

*Equal contribution

Running Head: Selection in genetic time series

Keywords:

Corresponding author:

Joshua B. Plotkin

Department of Biology

University of Pennsylvania

Philadelphia

PA 19104

USA

Phone

Fax

Email: jplotkin@sas.upenn.edu

Abstract

We develop a rigorous test for natural selection based on allele frequencies sampled from a population over multiple time points. We demonstrate that the standard method of estimating selection coefficients in this setting, and the associated chi-squared likelihood-ratio test of neutrality, is biased and it therefore does not provide a reliable test of selection. We introduce two methods to correct this bias, and we demonstrate that the new methods have power to detect selection in practical parameter regimes, such as those encountered in fitness assays of microbial populations. Our analysis is limited to a single diallelic locus, assumed independent of all other loci in a genome, which is again relevant to simple competition assays of laboratory and natural isolates; other techniques will be required to detect selection in time series of co-segregating, linked loci.

1 Introduction

Population geneticists typically seek to understand the forces responsible for patterns observed in contemporaneous samples of genetic data, such as the nucleotide differences fixed between species, polymorphisms within populations, and the structure of linkage disequilibrium. Recently, however, there has been a sharp increase in the abundance of dynamic data, where the frequencies of segregating alleles in an evolving population are monitored through time, both in laboratory experiments (HEGRENESS *et al.*, 2006; BARRICK *et al.*, 2009; LANG *et al.*, 2011) and in natural populations (BARRETT *et al.*, 2008; REID *et al.*, 2011; DENEFF and BANFIELD, 2012; WINTERS *et al.*, 2012; DANIELS *et al.*, 2012; PENNINGS *et al.*, 2012). One of the important questions is whether the changes in allele frequencies observed in such data are the result of natural selection or are simply consequences of genetic drift or sampling artifacts. In principle, it would seem that dynamic data might provide researchers with more power to detect and quantify selective forces while avoiding the assumptions of stationarity that are required for many inference techniques based on static samples (SAWYER and HARTL, 1992; DESAI and PLOTKIN, 2008; BOYKO *et al.*, 2008). Nonetheless, the behavior and power of inference techniques based on time series of genetic data have not been thoroughly investigated.

There is a well-developed literature on inferring population sizes from genetic time-series data assuming neutrality (POLLAK, 1983; WAPLES, 1989; WILLIAMSON and SLATKIN, 1999; WANG, 2001), but a much smaller literature on inferring natural selection from such time series (BOLLBACK *et al.*, 2008; ILLINGWORTH and MUSTONEN, 2011; ILLINGWORTH *et al.*, 2012; MALASPINAS *et al.*, 2012; MATHIESON and MCVEAN, 2013). Even the simplest case, which considers the dynamics of two alternative alleles at a single genetic locus, independent of all other loci, presents a number of statistical challenges that have not yet been addressed. BOLLBACK *et al.* (2008) fit two nested Wright-Fisher models to time-series data at a single locus (one model with selection, and one without), and rejected the neutral model using the χ^2 distribution for the likelihood ratio statistic. Such an approach is generally the most powerful and unbiased, at least for large datasets. Nonetheless, here we show that in practice the actual frequency of false positives under this approach can vastly exceed the nominal P -value obtained from the χ^2 distribution – and especially so at more stringent P -value cutoffs. Since the χ^2 distribution does not provide an accurate representation of the false positive rate, this approach cannot be used to draw sound statistical conclusions about the presence of selection from such time series. The underlying reason for this problem is that the likelihood ratio statistic is χ^2 -distributed only asymptotically, and convergence to this distribution is slow (WILKS, 1938). In most practical applications, such as when sampling from natural populations (REID *et al.*, 2011; DENEFF and BANFIELD, 2012; WINTERS *et al.*, 2012; DANIELS *et al.*, 2012; PENNINGS *et al.*, 2012) or competing two microbial strains (LENSKI *et al.*, 1991; BOLLBACK and HUELSENBECK, 2007), the number of sampled time points is small (less than 10), and the distribution of the likelihood ratio statistic is far from χ^2 under neutrality, leading to more false positives than expected.

We propose two solutions to fix this problem, providing unbiased tests for natural selection in time-series data sam-

pled at a single genetic locus. First, we develop an algorithm for computing the exact distribution of the likelihood ratio statistic under neutrality. Although feasible in many regimes, this direct approach is computationally intensive. We also propose an alternative, computationally efficient, albeit approximate, statistical method for rejecting the neutral model. Our approach builds directly on the work of BOLLBACK *et al.* (2008), and it is likewise limited to studying time series of allele frequencies at a single haploid locus, assuming independence from all other loci. The far more complicated problem of detecting selection from genomic time series of many linked loci has received attention elsewhere (ILLINGWORTH and MUSTONEN, 2011), and the problems identified here likely apply to those situations as well.

We start our presentation by introducing a likelihood framework for time-series data at a single genetic locus, and we then derive approximate expressions for these likelihoods under the Gaussian approximation to the Wright-Fisher model. We then demonstrate that the P -value given by the χ^2 distribution for the likelihood ratio statistic underestimates the actual false-discovery rate. Finally, we introduce two methods to correct this bias and we verify that they are virtually unbiased, even for small sample sizes. We conclude by quantifying the power of these methods to detect selection in different parameter regimes, considering also sampling noise in the allele frequency measurements.

2 Materials and Methods

2.1 General expression for the likelihood of time series allele-frequency data

Under any standard single-locus population-genetic model, the dynamics of an allele in a population are described by a Markov process that specifies the transition probability $P_s(x, t|x', t')$ that the allele assumes frequency x at time t , given that its frequency was x' at some previous time point t' . This transition probability depends on two parameters of interest, the population size N and the selection coefficient s acting on the allele (EWENS, 2004) and possibly also on other nuisance parameters. Ignoring sampling noise, the likelihood of observing allele frequencies $\nu_0, \nu_1, \dots, \nu_L$ at times $0, t_1, \dots, t_L$ is

$$L(\text{Data}; N, s) = U(\nu_0) \prod_{i=1}^L P_s(\nu_i, t_i | \nu_{i-1}, t_{i-1}), \quad (1)$$

and, under the neutral null hypothesis,

$$L(\text{Data}; N, 0) = U(\nu_0) \prod_{i=1}^L P_0(\nu_i, t_i | \nu_{i-1}, t_{i-1}). \quad (2)$$

Here $U(x)$ denote the probability of observing allele frequency x at time point 0 which for simplicity we set to be uniform on the interval $(0, 1)$, i.e., $U(x) \equiv 1$.

Given an algorithm for computing expressions (1) and (2), we can find the parameters values \hat{N} and \hat{s} that maximize the likelihood function (1) and the value \tilde{N} that maximizes the likelihood function (2) under the neutral null hypothesis.

We perform these maximization procedures using the golden section algorithm (KIEFER, 1953) implemented in the Gnu Scientific Library (GSL) package.

2.2 Computing the likelihood of time series data

Under most population-genetic models it is impossible to exactly compute the likelihood expressions (1), (2) because no exact analytic expressions for the transition probability $P_s(x, t|x', t')$ are available for arbitrary x, x', t, t' . Thus, various approximations have been proposed. The most well-known approximation to the Wright-Fisher and Moran models is the diffusion approximation of Kimura and others (EWENS, 2004). Although considerably simpler than the discrete Wright-Fisher or Moran models, the diffusion equation is still too difficult to solve exactly and efficiently in a general case. Although some methods are available (KIMURA, 1955a,b; EVANS *et al.*, 2007; BOLLBACK *et al.*, 2008; SONG and STEINRÜCKEN, 2012), they are rather cumbersome to implement numerically.

We therefore resort to the Gaussian approximation which is less general than the diffusion approximation but allows us to compute the transition probability efficiently, and quite accurately provided the allele has not been lost or fixed during the timescale of observation. We describe this approximation in detail in Appendix A. Briefly, if the time scale of observation is short compared to N , i.e., if absorption events can be ignored, the Moran process can be approximated by a sum of a deterministic process g and a Gaussian noise process Z (POLLETT, 1990). In the absence of genetic drift, i.e., when $N \rightarrow \infty$, the allele frequency X behaves deterministically, $X \rightarrow g(t, x_0)$, where g satisfies the logistic equation

$$\dot{g} = sg(1 - g), \quad (3)$$

$$g(0, x_0) = x_0, \quad (4)$$

whose solution is

$$g(t, x_0) = x_0 (x_0 + (1 - x_0)e^{-st})^{-1}. \quad (5)$$

Here x_0 is the initial deterministic allele frequency. When $N < \infty$, genetic drift perturbs the allele frequency X from its deterministic value and so $X(t) = g(t, x_0) + Z(t)$ where $Z(t)$ is the noise process. Then for any two time points $t > 0$ and $t' \geq 0$, the transition probability is approximated by

$$P_s(x, t|x', t') \approx \sqrt{\frac{N}{2\pi\sigma^2(\Delta t, g')}} \exp \left\{ -N \frac{(x - g - (x' - g')M(\Delta t, g'))^2}{2\sigma^2(\Delta t, g')} \right\}, \quad (6)$$

where

$$M(\Delta t, \xi) = e^{-s\Delta t} (\xi + (1 - \xi)e^{-s\Delta t})^{-2}, \quad (7)$$

$$\begin{aligned} \sigma^2(\Delta t, \xi) &= M^2(\Delta t, \xi)(2 + s)\xi(1 - \xi)s^{-1} \\ &\times \left[2\xi s(1 - \xi)\Delta t + \xi^2 e^{s\Delta t} - (1 - \xi)^2 e^{-s\Delta t} + ((1 - \xi)^2 - \xi^2) \right] \end{aligned} \quad (8)$$

and we used shorthands $g \equiv g(t, x_0)$, $g' \equiv g(t', x_0)$, $\Delta t = t - t'$. Under the neutral null hypothesis (i.e., when $s = 0$), the transition probability simplifies to

$$P_0(x, t|x', t') \approx \sqrt{\frac{N}{2\pi\sigma_n^2(\Delta t, x_0)}} \exp \left\{ -N \frac{(x - x')^2}{2\sigma_n^2(\Delta t, x_0)} \right\}, \quad (9)$$

with

$$\sigma_n^2(\Delta t, x_0) = 2x_0(1 - x_0)\Delta t. \quad (10)$$

Note that functions (5)–(10) depend on three parameters, N , s and x_0 . x_0 is a nuisance parameter that in principle can be fitted along with N and s . However, for the sake of reducing the number of fitted parameters, we fix x_0 to be equal to the observed allele frequency at time zero, $x_0 \equiv \nu_0$.

We assumed here and for the remainder of the paper that time is measured in generations (EWENS, 2004). If we measure time in physical units, equations (5)–(10) still hold, albeit with rescaled parameters $N \rightarrow N\tau$, and $s \rightarrow s/\tau$, where τ is the generation time.

3 Results & Discussion

We consider the problem of determining whether selection has played a role in shaping the fluctuations in the observed frequencies of an allele in a population over time. Suppose that at time points $0 < t_1 < \dots < t_L$ we sample n_0, n_1, \dots, n_L individuals from a given population of an unknown size N and find that b_i of them carry allele A at the locus of interest. Thus, we observe sampled allele frequencies $\nu_0 = b_0/n_0, \nu_1 = b_1/n_1, \dots, \nu_L = b_L/n_L$. We are concerned with the question of whether genetic drift and sampling alone are sufficient to explain the fluctuations in the observed frequency of allele A , or whether these frequency changes arose due to the action of natural selection either at the specified or another, completely linked, locus. Specifically, we consider the following pair of nested hypotheses. Under the neutral null hypothesis, changes in allele frequency are caused only by genetic drift and sampling, i.e., the selection coefficient s acting on allele A is fixed to be zero. Under the alternative hypothesis, there are no restrictions on s .

We initially ignore sampling noise. In other words, we assume that $n_i \gg 1$, $1 \ll b_i \ll n_i$ for all i , so that the

sampled allele frequencies ν_i accurately represent the actual frequencies in the population. We later investigate how sampling noise affects our conclusions.

3.1 Likelihood ratio statistic is not χ^2 distributed for finite data

Typically, testing the neutral null hypothesis involves constructing the likelihood ratio statistic (LRS) defined by

$$R(\text{Data}) = 2 \log \left(\frac{L(\text{Data}; \hat{N}, \hat{s})}{L(\text{Data}; \check{N}, 0)} \right), \quad (11)$$

where $L(\text{Data}; N, s)$ and $L(\text{Data}; N, 0)$ are the likelihoods of the data under the alternative and the null hypotheses, respectively; \hat{N} and \hat{s} maximize $L(\text{Data}; N, s)$; and \check{N} maximizes $L(\text{Data}; N, 0)$ (see Materials and Methods for details). The Neyman-Pearson lemma guarantees that the LRS defines the most powerful test of a given size (STUART *et al.*, 2009). In other words, the test that rejects the null hypothesis whenever $R(\text{Data}) > \kappa_\alpha$, with κ_α chosen so that the probability of a Type I error is α , has the lowest probability of Type II error among all tests that have the same probability of Type I error, α .

However, the Neyman-Pearson lemma holds only for *simple null hypotheses*, i.e., when the distributions of the observed random variables do not depend on parameters. In our problem, the null hypotheses is *composite* since the distributions of allele frequencies depend on an additional parameter, N , whose value is unknown. Importantly, this implies that the distribution of the likelihood ratio statistic is unspecified under the null hypothesis. Thus, not only is the likelihood ratio test not guaranteed to be the most powerful, there is no general way of determining the critical regions for the LRS distribution. The standard way to circumvent the latter problem is to use the asymptotic distribution for the LRS. When the number of data points approaches infinity, the LRS distribution converges to the χ^2 distribution with one degree of freedom, under appropriate regularity assumptions (WILKS, 1938). This approach has been previously used in the context of allelic time series by BOLLBACK *et al.* (2008).

Although the LRS is guaranteed to be asymptotically χ^2 distributed, the rate of convergence to this distribution is $O(L^{-1/2})$ where L is the number of sampled time points (WILKS, 1938). We set out to determine how well the χ^2 distribution approximates the true distribution of LRS when the number of data points is finite. This question is important because the use of an incorrect null distribution could result in a test that underestimates the fraction of Type I errors and thus erroneously rejects the null hypothesis more often than indicated by its P -value.

With this goal in mind, we simulated the neutral two-allele Wright-Fisher model with population size N , without mutation, with allele A initiated at 50%. We recorded the allele frequency every generation for T generations. We then produced a dataset consisting of these frequencies sampled every Δ generations. We considered a total of $L + 1$ sampled time points, so that $\Delta = T/L$. In order that the Gaussian approximation used to compute likelihoods is accurate, we made sure that absorption events are rare by setting $T \leq N/10$. For each population size N , we simulated

N	T	L	Δ	\tilde{N}/N	α			
					5%	1%	0.1%	0.01%
10^4	10	10	10^0	1.25	1.41	1.74	2.31	3.25
10^4	100	10	10^1	1.26	1.44	1.78	2.40	3.43
10^4	1000	10	10^2	1.29	1.64	2.37	4.26	8.08
10^3	100	10	10^1	1.29	1.70	2.49	4.54	8.26
10^4	1000	10	10^2	1.29	1.64	2.37	4.26	8.08
10^5	10000	10	10^3	1.31	1.23	1.54	2.26	3.58
10^4	100	5	20	1.69	2.25	3.67	7.29	13.93
10^4	100	10	10	1.29	1.70	2.49	4.54	8.26
10^4	100	100	1	1.07	1.32	1.84	3.73	9.66

Table 1: We report the ratio of the observed fraction of false positives α_{obs} in the likelihood ratio test to the fraction α expected under the assumption that LRS is χ^2 -distribution with 1 df, across a range of parameters (rows) and a range α values (columns 6–9). The P -values α expected under the χ^2 distribution are indicated in the top row. In the fifth column we also report the ratio of the population size most likely under the neutral null hypothesis \tilde{N} to its true value of N . See text for other notations.

10^4 allele-frequency trajectories. For each simulated allele-frequency time series we computed the LRS as described in Materials and Methods. Thus, for each combination of N , L and T we obtained the true distribution of LRS under the neutral null hypothesis. We calculated the critical values of LRS for a range of P -values, α , assuming that the LRS is distributed as χ^2 with one degree of freedom (df) and obtained the fraction of resulting false positives α_{obs} for each expected P -value.

The results shown in Table 1, Figure 1, and Supplementary Figure 1 demonstrate that the χ^2 distribution is a poor approximation for the true distribution of LRS under neutrality, when the number of sampled time points is finite. In particular, the P -value calculated from the χ^2 distribution with 1 df systematically underestimates the number of false positives α_{obs} even with as many as 100 sampled time points. While this discrepancy is moderate (less than a factor of 2) for relatively high nominal P -values (e.g., above 1%) it becomes increasingly more severe for more stringent P -values, so that in some regimes the χ^2 test will reject neutrality almost 10-times as often as it should (see Table 1).

As the third section of Table 1 shows, when the number of sampled points increases (and Δ decreases), with all other parameters held constant, the χ^2 P -value approaches an accurate representation of the true rate of false positive – that is, the ratios shown in the table approach 1.0. This result reflects the (slow) convergence of the LRS distribution to the χ^2 distribution with increasing amount of data, as guaranteed by the classical result of WILKS (1938).

In addition to the deviation of the true LRS distribution from the χ^2 distribution, the most likely value of N under the null hypothesis, \tilde{N} , systematically overestimates true N , especially when the number of data points is small (see Table 1). This upward bias in estimating population sizes has been previously reported by several authors (WAPLES, 1989; WILLIAMSON and SLATKIN, 1999; WANG, 2001). As expected, the bias in the estimation of N decreases as the number of data points increases, independent of the true population size or the number of generations represented in the sample (Supplementary Figure 1).

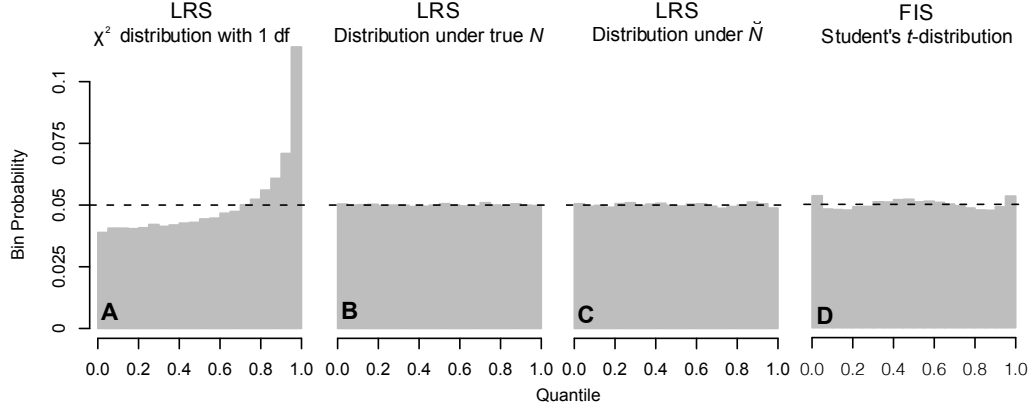


Figure 1: Comparison of distributions of test statistics with their approximations. Shown is the probability that a test statistic falls within each quantile of a given distribution. **Panels A–C.** The probability distribution for LRS to fall into each quantile of (A) the χ^2 distribution with 1 df; (B) the distribution under the true N ; (C) the empirical LRS distribution under \tilde{N} . The LRS statistic falls into the top quantiles of the χ^2 distribution with 1 df more often than expected which indicates that the P -value given by the χ^2 distribution underestimates the true probability of a Type I error. The distribution of LRS under \tilde{N} closely approximates the LRS distribution under true N . **Panel D.** The probability distribution for the FIS to fall into each quantile of the Student’s t -distribution with L df. Student’s t distribution is a good approximation for the true distribution of the t -statistic. Parameter values used: $N = 10^3$, $\Delta = 20$, $L = 5$. We ran 200,000 and 300,000 simulations to estimate the LRS distribution under true N and under \tilde{N} , respectively.

3.2 “Empirical” likelihood ratio test

In order to fix the shortcomings of the χ^2 test for selection we introduced two new methods. The most direct approach would be to compare the LRS statistic observed in the data to the true distribution of the LRS statistic assuming neutrality. However, this approach is not possible, because in practice we do not know the true value of the population size N under which to generate the distribution of the LRS. Therefore, as an intuitive approximation, we propose first to estimate the most-likely population size assuming neutrality, \tilde{N} , to generate, via neutral Wright-Fisher simulations, the empirical LRS distribution with \tilde{N} as the population size, and to compare the LRS observed in the data against this empirical LRS distribution to determine whether or not to reject neutrality. We call this approach the empirical likelihood ratio test (ELRT).

Figure 1 shows that the empirical LRS distribution is an excellent approximation to the true LRS distribution even when the number of samples is small. Although the P -values computed with the empirical LRS distribution are accurate, this approach is relatively computationally intensive because many Wright-Fisher simulations must be performed to obtain the empirical distribution, and each simulation must be accompanied by the calculation of the LRS. To reduce the computational burden, in the next section we propose an alternative to the likelihood ratio test that is computationally extremely easy, but is somewhat less accurate.

3.3 “Frequency increment” test

We define the “frequency increment” (FI) statistic as

$$t_{\text{FI}}(\text{Data}) = \frac{\bar{Y}}{\sqrt{S^2/L}}, \quad (12)$$

where \bar{Y} and S are the sample mean and the sample variance

$$\bar{Y} = \frac{1}{L} \sum_{i=1}^L Y_i \quad \text{and} \quad S^2 = \frac{1}{L-1} \sum_{i=1}^L (Y_i - \bar{Y})^2$$

of the rescaled allele frequency increments

$$Y_i = \frac{\nu_i - \nu_{i-1}}{\sqrt{2\nu_{i-1}(1 - \nu_{i-1})(t_i - t_{i-1})}}, \quad i = 1, 2, \dots, L.$$

Since under the neutral null hypothesis the allele frequency ν behaves approximately as Brownian motion as long as it is far from the boundaries 0 or 1 (see EWENS (2004) and Appendix), the random variables Y_i are independent and approximately normally distributed with mean 0 and variance 1 (see equation (10)). Thus, the distribution of the FI statistic under neutrality can be approximated by the Student’s t -distribution with L degrees of freedom. We call this test the frequency increment test (FIT).

Figure 1 and Table 2 show that the test of neutrality based upon the FI statistic coupled with the Student’s t -distribution substantially outperforms the likelihood ratio test with the χ^2 distribution – in the sense that it is far less biased. Nevertheless, the nominal P -value based on the t -distribution with L df does not always accurately reflect the true probability of Type I error. Under most parameter regimes where the probability of Type I error deviates from the P -value reported by the t -distribution, the FI test is overly conservative (i.e., the reported P -value overestimates the probability of Type I error), but the precise dependence involves N , L , Δ in a complicated way (Table 2). In any case, the deviation in the size of this test is an order of magnitude smaller than that of the LRT, in all parameter regimes tested.

3.4 Power of ELRT and FIT to detect selection

Next we tested whether the empirical likelihood ratio test (ELRT) and the frequency increment test (FIT) have power to detect selection in allele-frequency data. To this end, we ran Wright-Fisher simulations as described above but now with one of the two alleles having selective advantage s . For each N and s we simulated 10^4 allele-frequency trajectories. As before, we sampled $L + 1$ frequencies of the selected allele spaced Δ generations apart. For each combination of N , s , L and Δ we performed ELRT and FIT, and we calculated how often each test reports a significant outcome for a

N	T	$L + 1$	Δ	α			
				5%	1%	0.1%	0.01%
10^4	8	5	2	1.00	1.02	1.03	1.10
10^4	98	50	2	0.99	0.97	0.93	0.90
10^4	998	500	2	0.83	0.69	0.46	0.28
10^4	9	10	10^0	1.00	1.01	0.97	1.01
10^4	90	10	10^1	0.99	1.00	1.01	1.06
10^4	900	10	10^2	0.96	1.05	1.22	1.24
10^3	90	10	10^1	0.99	1.10	1.34	1.47
10^4	900	10	10^2	0.96	1.05	1.22	1.24
10^5	9000	10	10^3	0.53	0.42	0.21	0.12

Table 2: We report the ratio of the observed fraction of false positives α_{obs} in the FI test to the fraction α expected under the assumption that the FI statistic (12) is distributed according to Student’s t -distribution with L df, across a range of parameters (rows) and a range α values (columns 5–8). The P -values α are indicated in the top row. See text for other notations.

given P -value cutoff.

As shown in Figure 2, both the ELRT and the FIT have substantial power to detect the action of natural selection when selection is moderate to strong ($Ns > 10$), with FIT having slightly more power than the ELRT. The power of both tests strongly depends on the amount of time that the allele frequency is observed: longer observations yield substantially more power. The power to detect selection depends only weakly on the number of sampled time points and on the population size, all other parameters being equal.

3.5 The effect of sampling

So far we demonstrated that we can accurately calculate the sizes of the empirical likelihood ratio test and the frequency increment test (FIT) and that these tests have substantial power to detect selection in an idealized situation when the exact allele frequency in the population is known. We now investigate the behavior of these tests in a more realistic situation when the allele frequency is measured by taking a sampling and typing a finite number of individuals from the population. We used the same time-series trajectories as in previous sections, but now instead of analyzing the true allele frequencies x we drew binomial random variables with sample size n and success probability x to obtain the sampled allele frequencies ν . We then analyzed the test size and the test power using sampled allele frequencies ν as the data. Figure 3 shows that the P -values estimates for both ELR and FI tests remain unbiased when the allele frequencies are sampled. However, as expected, the power of both tests to reject neutrality declines as the number of samples decreases (Figure 4). If the number of samples falls below 100 per sampled time point, i.e., when sampling noise becomes strong, we lose the ability to detect selection unless the selection coefficient is very high (e.g., $Ns > 20$) or the time series is very long.

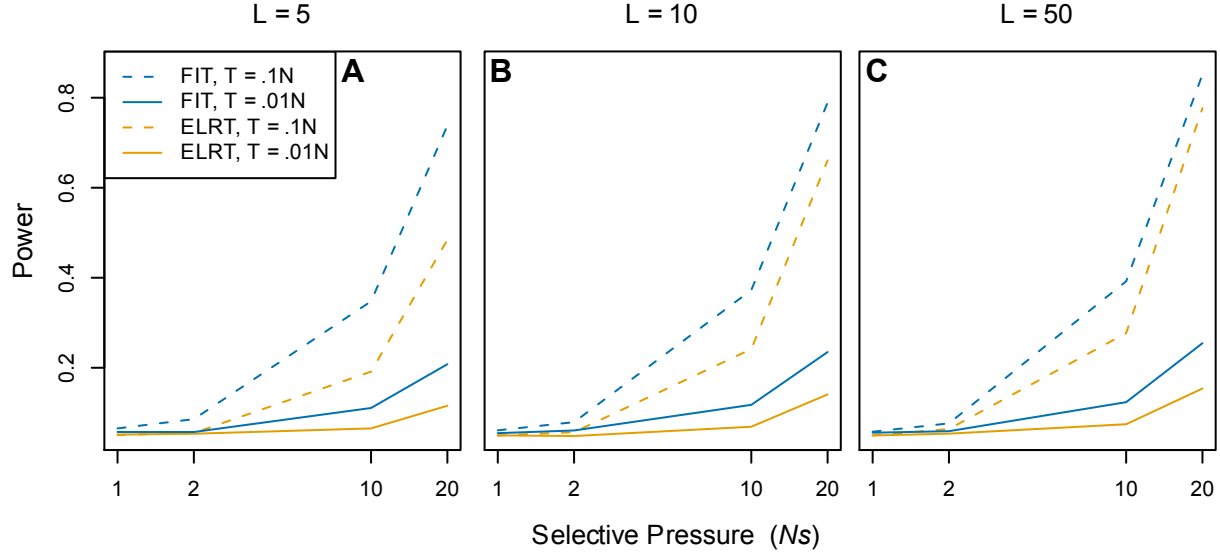


Figure 2: Power of ELRT and FIT as a function of strength of selection. Power is the fraction of trail data sets generated by the Wright-Fisher model with selection for which ELRT (yellow lines) or FIT (blue lines) rejects the neutral null hypothesis at P -value $\alpha = 0.05$. Dashed (solid) lines indicate data sets where allele frequency was sampled for $T = N/10$ ($T = N/100$) generations. Each point represents an average over 10^3 trials, with $N = 10^4$.

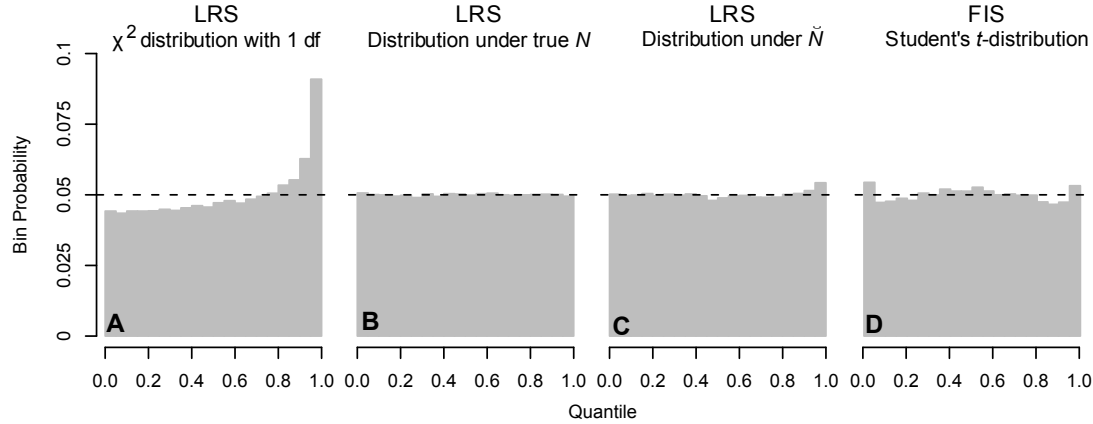


Figure 3: Comparison of distributions of test statistics with their approximations, allele frequencies are sampled Shown is the probability that a test statistic falls within each quantile of a given distribution when allele frequencies are sampled. **Panels A–C.** The probability distribution for LRS to fall into each quantile of (A) the distribution under the true N ; (B) the χ^2 distribution with 1 df; (C) the LRS distribution under \tilde{N} . **Panel D.** The probability distribution for the FIS to fall into each quantile of the Student's t -distribution with L df. Student's t distribution is a good approximation for the true distribution of the t -statistic. Parameter values used: $N = 1000$, $\Delta = 20$, $L = 5$, $n = 500$. We ran 200,000 and 300,000 simulations to estimate the LRS distribution under true N and under \tilde{N} , respectively.

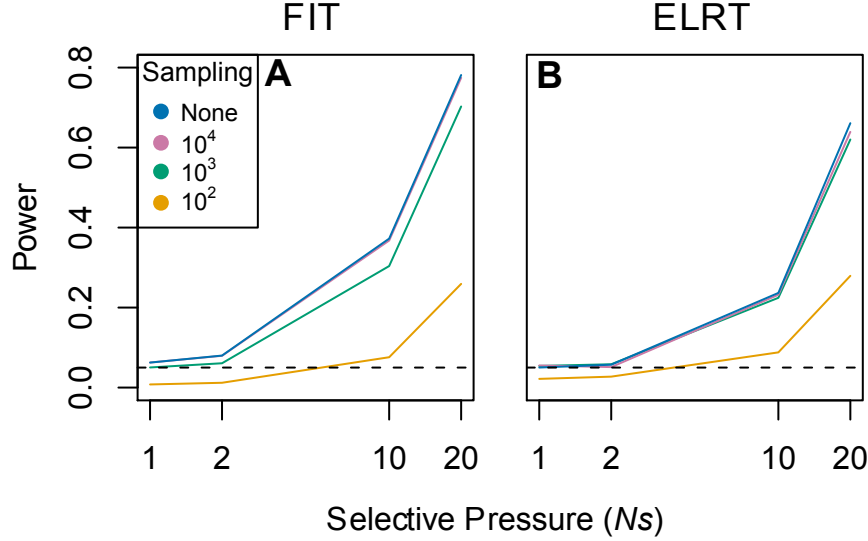


Figure 4: Power of ELRT and FIT as a function of strength of selection under various sampling regimes. Each point represents an average over 10^3 trials, except for ELRT with 10^4 subsamples, which used only 2,500 trials.

4 Conclusions

We have demonstrated that the standard, χ^2 -based test for selection in time series of allele frequencies (BOLLBACK *et al.*, 2008) is subject to a greatly elevated false discovery rate, in the practical regime of relatively few sampled time points. As a result of this bias, the LRT is not reliable for testing selection in many practical time series, because its P -value underestimates the rate of false inferences of selection. We have presented two methods to address this problem, developing likelihood-based tests for selection that avoid the χ^2 assumption. We have shown that both our methods are unbiased and have power to detect selection in parameter regimes that are reasonable for many microbial experiments and natural populations.

Specifically, our results indicate that our proposed methods should be able to detect selection in the ranges relevant for many microbial and viral populations – that is $N > 10^4$, $s > 0.001$, and possibly even smaller selection coefficients in larger populations. These techniques, just as the χ^2 based techniques, will be limited to the simple situations in which the frequencies of alleles observed at a locus are not influenced by mutations that possibly arise elsewhere in the genome, during the time of observation. This simple situation, in which clonal interference is absent, applies to the competitive fitness assays used in evolution experiments (LENSKI *et al.*, 1991; GALLET *et al.*, 2012) or to tracking known polymorphisms in natural populations for a relatively short time (BARRETT *et al.*, 2008; WINTERS *et al.*, 2012; DANIELS *et al.*, 2012; PENNINGS *et al.*, 2012). By contrast, inferring selection coefficients when allele dynamics are influenced by multiple linked sites is a substantially more difficult problem, which has begun to be addressed elsewhere (ILLINGWORTH and MUSTONEN, 2011; ILLINGWORTH *et al.*, 2012), although not, of course, with an rigorous likelihood framework for the dynamics of all possible genotypes.

5 Appendix A. Gaussian approximation to the Moran process

We approximate the continuous-time Moran processes with a combination of a deterministic process and Gaussian noise process. We follow here the procedure outlined by POLLETT (1990). The Moran's stochastic process describes the number $n^{(N)}(t)$ of mutants in a population of constant size N at time t . This number can increase by one from i to $i + 1$ with rate

$$r^{(N)}(i, i + 1) = \mu_m i \frac{\lambda_w(N - i)}{\lambda_w(N - i) + \lambda_m i}$$

and decrease by one with rate

$$r^{(N)}(i, i - 1) = \mu_w(N - i) \frac{\lambda_m i}{\lambda_w(N - i) + \lambda_m i}.$$

Here, μ_w and λ_w are the birth and death rates of the wildtype, and μ_m and λ_m are the birth and death rates of the mutant type, respectively. We assume $\lambda_w = \lambda_m$, $\mu_w = 1$, and let $\mu_m = (1 + s)\mu_w = 1 + s$. Then

$$r^{(N)}(i, i + 1) = N f_{+1}(i/N), \quad r^{(N)}(i, i - 1) = N f_{-1}(i/N) \quad (13)$$

with

$$f_{+1}(x) = (1 + s)x(1 - x), \quad f_{-1}(x) = x(1 - x).$$

Define

$$\begin{aligned} F(x) &= \sum_{\delta \in \{-1, +1\}} \delta f_{\delta}(x) = sx(1 - x) \\ G(x) &= \sum_{\delta \in \{-1, +1\}} \delta^2 f_{\delta}(x) = (2 + s)x(1 - x). \end{aligned}$$

Let $X^{(N)}(t) = n^{(N)}(t)/N$ be the frequency of the mutant in the population at time t . The limit of $X^{(N)}$, $g(t, x_0) = \lim_{N \rightarrow \infty} X^{(N)}(t)$, is a deterministic function that, under certain regularity conditions, satisfies equations (3), (4) with $x_0 = \lim_{N \rightarrow \infty} X^{(N)}(0)$ and solution given by (5).

Now let

$$Z(t) = \lim_{N \rightarrow \infty} \sqrt{N} \left(X^{(N)}(t) - g(t, x_0) \right) \quad (14)$$

be the asymptotic process that describes the noise around the deterministic trajectory. If we knew the distribution of $Z(t)$, we could approximate the frequency $X^{(N)}$ at a finite N by

$$X^{(N)}(t) \approx g(t, x_0) + \frac{1}{\sqrt{N}} Z(t). \quad (15)$$

The asymptotic noise process is in general a diffusion process, but, as long as it remains far from absorbing boundaries, it can be approximated by a Gaussian process with the corresponding first two moments. The advantage of this approach is that the first two moments of the diffusion process can be computed analytically, resulting in an expression for the probability distribution of the allele frequency at time t .

If $z_0 = \lim_{N \rightarrow \infty} \sqrt{N}(X^{(N)}(0) - x_0)$ is the initial value of the limiting noise process, then the mean and variance of the noise process at time $t \geq 0$ are $\mathbb{E}Z(t) = M(t, x_0)z_0$ and $\text{Var } Z(t) = \sigma^2(t, x_0)$ respectively, where $M(t, x_0)$ satisfies the equations

$$\frac{dM}{dt} = F'(g(t, x_0))M = s \frac{(1 - x_0)e^{-st} - x_0}{(1 - x_0)e^{-st} + x_0} M \quad (16)$$

$$M(0, x_0) = 1 \quad (17)$$

and $\sigma^2(t, x_0)$ satisfies the equations

$$\begin{aligned} \frac{d\sigma^2}{dt} &= 2F'(g(t, x_0))\sigma^2 + G(g(t, x_0)) = \\ &= \frac{(1 - x_0)e^{-st} - x_0}{(1 - x_0)e^{-st} + x_0} 2s\sigma^2 + \frac{(2 + s)x_0(1 - x_0)e^{-st}}{((1 - x_0)e^{-st} + x_0)^2} \end{aligned} \quad (18)$$

$$\sigma^2(0, x_0) = 0. \quad (19)$$

The solution to equations (16), (17) is given by

$$M(t, x_0) = \exp \left\{ \int_0^t F'(g(\tau, x_0)) d\tau \right\},$$

that, after substituting F' and g , yields

$$M(t, x_0) = e^{-st} (x_0 + (1 - x_0)e^{-st})^{-2}.$$

The solution to equations (18), (19) is given by

$$\sigma^2(t, x_0) = M^2(t, x_0) \int_0^t M^{-2}(\tau, x_0) G(g(\tau, x_0)) d\tau,$$

that, after substituting G and g , yields

$$\begin{aligned} \sigma^2(t, x_0) &= M^2(t, x_0)(2 + s)x_0(1 - x_0)s^{-1} \\ &\times \left[2x_0s(1 - x_0)t + x_0^2e^{st} - (1 - x_0)^2e^{-st} + ((1 - x_0)^2 - x_0^2) \right]. \end{aligned}$$

If the true state of the stochastic process $X^{(N)}$ is known to be $X^{(N)}(0)$ at the time point 0, we can approximate the initial value of the limiting noise process as $z_0 \approx \sqrt{N}(X^{(N)}(0) - x_0)$. Then from (15) we have

$$\begin{aligned}\mathbb{E}X^{(N)}(t) &\approx g(t, x_0) + M(t, x_0) \left(X^{(N)}(0) - x_0 \right), \\ \text{Var } X^{(N)}(t) &\approx \frac{1}{N} \sigma^2(t, x_0).\end{aligned}$$

Analogously, if the value of the process $X^{(N)}$ is known to be $X^{(N)}(t')$ at a later time $t' \geq x_0$, then at time $t \geq t'$ we have

$$\mathbb{E}_{t'} X^{(N)}(t) \approx g(t, x_0) + M(\Delta t, g(t', x_0)) \left(X^{(N)}(t') - g(t', x_0) \right), \quad (20)$$

$$\text{Var}_{t'} X^{(N)}(t) \approx \frac{1}{N} \sigma^2(\Delta t, g(t', x_0)), \quad (21)$$

where $\Delta t = t - t'$, and $\mathbb{E}_{t'}$ and $\text{Var}_{t'}$ denote conditional expectation and variance given the state of the process at time t' . Thus, the conditional distribution of the allele frequency $X^{(N)}$ at time t given its value at time $t' \leq t$ can be approximated by a Gaussian distribution with mean given by (20) and variance given by (21). We apply this approximation to every observation interval $(t_{i-1}, t_i), i = 1, \dots, L$. As noted above, the initial value of the deterministic process, x_0 , is a free parameter that can be fitted along with N and s . However, we set x_0 to be equal to the observed allele frequency ν_0 at time 0 in order to reduce the number of fitted parameters.

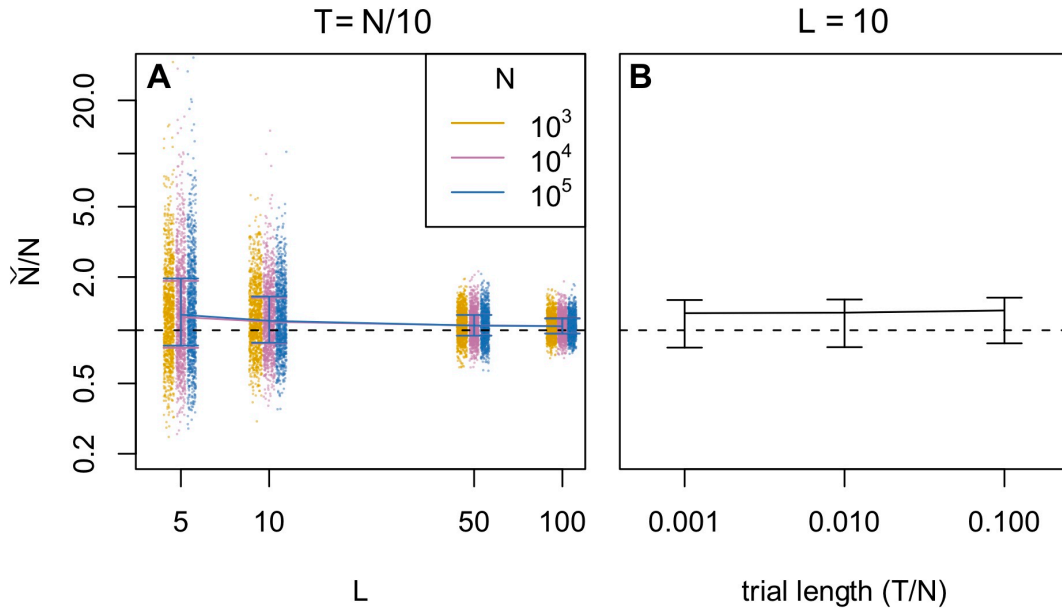
Note that the approximations described here work for the Moran process which is density-dependent as can be seen from equations (13). The Wright-Fisher process is not density-dependent and, strictly speaking, the approximations described here are not valid, although in practice they work well.

References

- BARRETT, R. D. H., S. M. ROGERS, and D. SCHLUTER, 2008 Natural selection on a major armor gene in threespine stickleback. *Science* **322**: 255–257.
- BARRICK, J. E., D. S. YU, S. H. YOON, H. JEONG, T. K. OH, *et al.*, 2009 Genome evolution and adaptation in a long-term experiment with *Escherichia coli*. *Nature* **461**: 1243–1247.
- BOLLBACK, J. P., and J. P. HUELSENBECK, 2007 Clonal interference is alleviated by high mutation rates in large populations. *Mol. Biol. Evol.* **24**: 1397–1406.
- BOLLBACK, J. P., T. L. YORK, and R. NIELSEN, 2008 Estimation of $2N_e s$ from temporal allele frequency data. *Genetics* **179**: 497–502.

- BOYKO, A. R., S. H. WILLIAMSON, A. R. INDAP, J. D. DEGENHARDT, R. D. HERNANDEZ, *et al.*, 2008 Assessing the evolutionary impact of amino acid mutations in the human genome. *PLoS Genet* **4**: e1000083.
- DANIELS, R., H.-H. CHANG, P. D. SÉNE, D. C. PARK, D. E. NEAFSEY, *et al.*, 2012 Genetic surveillance reveals shift to clonal transmission of malaria following enhanced intervention in Senegal. In preparation .
- DENEV, V. J., and J. F. BANFIELD, 2012 In situ evolutionary rate measurements show ecological success of recently emerged bacterial hybrids. *Science* **336**: 462–466.
- DESAI, M. M., and J. B. PLOTKIN, 2008 The polymorphism frequency spectrum of finitely many sites under selection. *Genetics* **180**: 2175–2191.
- EVANS, S. N., Y. SHVETS, and M. SLATKIN, 2007 Non-equilibrium theory of the allele frequency spectrum. *Theor. Pop. Biol.* **71**: 109–119.
- EWENS, W. J., 2004 *Mathematical population genetics..* Springer Science + Business Media, Inc., New York.
- GALLET, R., T. F. COOPER, S. F. ELENA, and T. LENORMAND, 2012 Measuring selection coefficients below 10^{-3} : Method, questions, and prospects. *Genetics* **190**: 175–186.
- HEGRENESS, M., N. SHORESH, D. HARTL, and R. KISHONY, 2006 An equivalence principle for the incorporation of favorable mutations in asexual populations. *Science* **311**: 1615–1617.
- ILLINGWORTH, C. J., L. PARTS, S. SCHIFFELS, G. LITI, and V. MUSTONEN, 2012 Quantifying selection acting on a complex trait using allele frequency time series data. *Mol. Biol. Evol.* **29**: 1187–1197.
- ILLINGWORTH, C. J. R., and V. MUSTONEN, 2011 Distinguishing driver and passenger mutations in an evolutionary history categorized by interference. *Genetics* **189**: 989–1000.
- KIEFER, J., 1953 Sequential minimax search for a maximum. *Proc. Am. Math. Soc.* **4**: 502–506.
- KIMURA, M., 1955a Solution of a process of random genetic drift with a continuous model. *Proc. Natl. Acad. Sci. USA* **41**: 144–150.
- KIMURA, M., 1955b Stochastic processes and distribution of gene frequencies under natural selection. *Cold Spring Harb Symp Quant Biol* **20**: 33–53.
- LANG, G. I., D. BOTSTEIN, and M. M. DESAI, 2011 Genetic variation and the fate of beneficial mutations in asexual populations. *Genetics* **188**: 647–661.
- LENSKI, R. E., M. R. ROSE, S. C. SIMPSON, and S. C. TADLER, 1991 Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. *Am Nat* **138**: 1315–1341.

- MALASPINAS, A. S., O. MALASPINAS, S. N. EVANS, and M. SLATKIN, 2012 Estimating allele age and selection coefficient from time-serial data. *Genetics* **192**: 599–607.
- MATHIESON, I., and G. MCVEAN, 2013 Estimating Selection Coefficients in Spatially Structured Populations from Time Series Data of Allele Frequencies. *Genetics* .
- PENNINGS, P., S. KRYAZHIMSKIY, and J. WAKELEY, 2012 Loss and recovery of genetic diversity in adapting populations of HIV. In preparation .
- POLLAK, E., 1983 A new method for estimating the effective population size from allele frequency changes. *Genetics* **104**: 531–548.
- POLLETT, P. K., 1990 On a model for interference between searching insect parasites. *J. Austral. Math. Soc. Ser. B* **32**: 133–150.
- REID, B. J., R. KOSTADINOV, and C. C. MALEY, 2011 New strategies in Barrett's esophagus: Integrating clonal evolutionary theory with clinical management. *Clin Cancer Res* **17**: 3512–3519.
- SAWYER, S. A., and D. L. HARTL, 1992 Population genetics of polymorphism and divergence. *Genetics* **132**: 1161–1176.
- SONG, Y. S., and M. STEINRÜCKEN, 2012 A simple method for finding explicit analytic transition densities of diffusion processes with general diploid selection. *Genetics* **190**: 1117–1129.
- STUART, A., K. ORD, and S. ARNOLD, 2009 *Kendall's advanced theory of statistics, classical inference and the linear model*. John Wiley & Sons.
- WANG, J., 2001 A pseudo-likelihood method for estimating effective population size from temporally spaced samples. *Genet Res* **78**: 243–257.
- WAPLES, R. S., 1989 A generalized approach for estimating effective population size from temporal changes in allele frequency. *Genetics* **121**: 379–391.
- WILKS, S. S., 1938 The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Stat.* **9**: 60–62.
- WILLIAMSON, E. G., and M. SLATKIN, 1999 Using maximum likelihood to estimate population size from temporal changes in allele frequencies. *Genetics* **152**: 755–761.
- WINTERS, M. A., R. M. LLOYD, R. W. SHAFER, M. J. KOZAL, M. D. MILLER, *et al.*, 2012 Development of elvitegravir resistance and linkage of integrase inhibitor mutations with protease and reverse transcriptase resistance mutations. *PLoS One* **7**: e40514.



Supplementary Figure 1: Bias in the maximum-likelihood estimate of population size, under neutrality. Shown is the ratio of the ML estimated population size \tilde{N} to the true population size N as a function of the number of sampled points L (left panel) and as a function of the length of the observed time series $L\Delta$ (right panel). Bars indicate ± 2 standard deviation of the distribution of \tilde{N}/N . As expected, bias in \tilde{N} decreases as the number of sampled time points increases. The bias is almost independent of N and the length of the sampling period. Parameters used: (A) 100,000 trials for each N , 1000 points displayed, $\Delta = 1$; (B) 100,000 for each L/N with $N = 10^4$, $\Delta = 1$.